




PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04L 12/44	A1	(11) International Publication Number: WO 99/29071 (43) International Publication Date: 10 June 1999 (10.06.99)
(21) International Application Number: PCT/GB98/03575 (22) International Filing Date: 30 November 1998 (30.11.98) (30) Priority Data: 9725368.6 28 November 1997 (28.11.97) GB (71) Applicants: 3COM TECHNOLOGIES [---]; Ugland House, P.O. Box 309, Georgetown, Grand Cayman (KY). BUTCHER, Ian, James [GB/GB]; A.A. Thornton & Co., Northumberland House, 303-306 High Holborn, London WC1V 7LE (GB). (72) Inventors: MURPHY, Ciaran; 6 Kempton Heath, Navan Road, Dublin 7 (IE). CLIFFORD, Neil; 9 Wasdale Park, Terenure, Dublin 6 (IE). NOLAN, Jerome; 8 Sandyford Downs, Sandyford, Dublin 18 (IE). GLEESON, Mick; 2 Royal Canal Court, Philbsborough, Dublin 7 (IE). (74) Common Representatives: BUTCHER, Ian, James et al.; A.A. Thornton & Co., Northumberland House, 303-306 High Holborn, London WC1V 7LE (GB).		(81) Designated States: GB, JP, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>

(54) Title: RESOURCE SHARING**(57) Abstract**

Efficient access to a shared resource such as buffer memory by a number of users is achieved by providing a single access control device which is supplied with parameters appropriate to the user having access to the resource at the relevant time. In particular, in a network communication device in which MAC devices have a common buffer memory, appropriate pointers related to the users are supplied to a calculation means during access to the buffer by the various users.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

Resource Sharing

The present invention relates to an arrangement for efficiently controlling access to a shared resource, such as a shared memory, by a number of users of the memory. This is described in the context of a network interface for a computer network communication device but it is applicable to other areas as well.

In computer networks, computers and other computing devices are interconnected by way of communication devices to enable them to communicate with each other. In particular, a communication device typically has a plurality of ports to each of which a computing device can be connected. The communication device receives communications via its ports and, according to a number of well known schemes, retransmits the received communications via the other ports to enable the communications to reach their intended destination or destinations. To enable both the computing devices and the communication devices to communicate via the network, the devices are provided with network interfaces, known as media access controllers (MACs) which provide the necessary buffering etc. for connecting onto wires forming the network. A single computing device having a single network connection will simply be provided with a single MAC.

However, a communication device which forms part of the structure of the network has, as mentioned above, a possibly large plurality of ports. There will be a MAC provided for each of these ports to enable proper interfacing on the network. In particular in the transmit side of a MAC, it is known to provide some buffering of the communications which are to be transmitted on to the network and therefore it may be considered that a typical transmit MAC has a front end, via which communications for transmission onto the network are received, and a back end, via which the communications are transmitted to the network. Between the front end and the back end there is typically provided a buffer memory in which

the communications are temporarily stored.

In a communication device as mentioned above where there are a plurality of MACs provided, the plurality of buffers which are therefore required by the plurality of the transmit MACs may be provided in a single physical memory with each transmit MAC being assigned a separate portion of the physical memory. Typically, there would also be arbitration means for allowing access to the physical memory at different times by different ones of the MACs. Such a shared memory resource is particularly useful used in conjunction with an arrangement in which a plurality of MACs are implemented on a single physical device and the memory may be a simple static RAM (SRAM) associated with the integrated circuit embodying the plurality of MACs.

Associated with each MAC it is necessary to store a number of parameters or pointers associated with the buffer memory used by that MAC. Typically, each buffer memory is simply a circular queue in a designated block of SRAM and in this case it is necessary to store an indication of the current start and finish of the occupied portion of the queue, together with other pointers necessary for the proper functioning of the MAC, as will be described in more detail below. It is also necessary to perform a number of arithmetic operations on the basis of the stored parameters and previously, in the case where a plurality of MACs are implemented on the same physical device, a separate computation means has been provided for each MAC, having associated therewith a plurality of registers storing the required parameters.

As the demand for the number of ports has increased it becomes impractical simply to replicate the computation device for each MAC, as this considerably increases the gate count, and therefore, costs, of the chip which implements the arrangement.

The present invention provides a communication device for a computer network in which a plurality of network devices are interconnected and enabled to send communications to each other, the device comprising:

a plurality of ports via which communications may be received and re-transmitted;

a communications core arranged to re-transmit received communications via one or more of said ports according to a defined functionality;

5 transmit means associated with each of said ports each arranged to receive from said communications core communications to be transmitted via the respective port and to transmit said communications via said respective port;

buffer memory means associated with said transmit means and arranged to have temporarily stored therein communications received by said transmit means prior to their transmission, each port having associated therewith
10 a plurality of pointers related to the corresponding use by the transmit means of said buffer memory means;

buffer memory access means arranged to control access to said buffer memory means by said transmit means for each of said plurality of ports
15 according to the corresponding plurality of pointers; and

pointer memory means arranged to store said pointers for each of said ports and to provide to said buffer memory access means said pointers associated with one of said ports when said buffer access means permits access to said buffer memory means by said transmit means for said one of said ports.

20 The present invention also provides a method for controlling access to a shared resource by a plurality of users of that resource, each of said users having associated therewith a plurality of parameters associated with said shared resource on the basis of which calculations need to be done in connection with that user's access to the resource, the method comprising storing all of the parameters
25 associated with said plurality of users, providing a single calculation means, and providing the parameters associated with each user to said calculation means to enable it to perform said calculations at times according to the times each user has access to the shared resource.

In a preferred arrangement, the corresponding parameters for all of

the users are stored together and simply indexed according to which user they relate to. This facilitates the arrangement of the single calculation device as the appropriate parameters can simply be read from all of the different storage locations according to which user is currently under consideration.

5 In the context of the above discussed multiple MAC situation for a communication device this invention is used in conjunction with the arbitration of access to the memory to enable calculations associated with each relevant MAC to be conducted at appropriate times. In an arrangement where twelve MACs are to be implemented on a single chip, the removal of the requirement for there to be
10 a separate calculation device associated with each MAC provides the advantage of reducing the number of gates on the chip by approximately fifty thousand.

 In order that this invention may be better understood a preferred embodiment will be described in the following in the context of the above mentioned multiple MAC arrangement, by way of example and with reference to
15 the accompanying figures in which:

 Figure 1 is a schematic illustration of a communication device;

 Figure 2 illustrates a schematic diagram of a known arrangement in which a plurality of MAC devices are implemented on a single chip; and

 Figure 3 illustrates a corresponding arrangement according to the
20 preferred embodiment of the present invention.

 In this invention efficient access to a shared resource such as buffer memory by a number of users is achieved by providing a single access control device which is supplied with parameters appropriate to the user having access to the resource at the relevant time. In particular, in a network communication
25 device in which MAC devices have a common buffer memory, appropriate pointers related to the users are supplied to a calculation means during access to the buffer by the various users.

 This invention is described below in the context of access to a buffer memory by MAC devices in a network communication device. The actual

allocation of the buffer memory to the various users is not described in detail herein, but may be conducted according to our invention described in co-pending Application _____ entitled "Dynamic Memory Allocation" and filed on even date herewith.

5 Figure 1 is a schematic illustration of a communication device 1 in which this invention may be implemented. Communication device 1 has a plurality of ports 2 via which network communications are received and transmitted, each of which has a media access control (MAC) device 6 associated with it. As is well known, MAC devices 6 control the operation of the respective
10 ports and the handling of the data by the respective ports. At the heart of communication device 1 is communication core means 4 which provides the basic functionality of device 1, eg. a repeater or bridge functionality, as well as, for instance, management and control functions. Data received at the ports 2 is passed to communications core 4 and data for re-transmission is passed to ports 2 via
15 appropriate interconnection means, for instance bus arrangement 12 illustrated in Figure 1.

 Also provided within communication device 1 is buffer memory 40, typically in the form of SRAM. Each of MAC devices 6 is able to communicate with buffer memory 40 in order to buffer data passing through the respective port.
20 This buffering is principally done for data which is being transmitted, and the following more detailed description is therefore in terms of the transmit side of MAC devices 6, by way of example.

 Figure 2 illustrates in more detail a known arrangement in which a plurality of MAC devices are implemented within device 1 on a single chip. The
25 multiple MAC device 10 comprises a plurality of MACs, of which only three are illustrated for clarity. The transmit side of each MAC device comprises a front end portion 20 and a back end portion 30. Communications to be transmitted by the plurality of MACs in the multiple MAC arrangement 10 are received, from data bus 12, via input 14 and are routed in a known fashion to the appropriate one

of the plurality of front end MAC portions 20 on chip 10. Associated with the multiple MAC arrangement is SRAM 40 which is used for buffering in the MAC arrangement. In particular, each MAC device on chip 10 has associated with it a specific portion of SRAM 40 in which front end portion 20 can store a received communication until back end portion 30 is ready to transmit it to the network. Access to SRAM 40 is controlled by SRAM arbiter 44 and is via multiplexer 42. The control of shared access to a memory resource is well known and will not be described in detail here.

Each MAC further comprises a calculation device 24 which performs the necessary calculations associated with the access by that MAC to the SRAM 40. In order to perform these calculations, it is necessary to know the values of certain parameters relating to the portion of SRAM assigned to that MAC and for this purpose, there are provided a plurality of registers 26 storing the current values of the appropriate parameters. Typically, each MAC's portion of memory will be arranged as a cyclic FIFO queue and in order for this to function properly it is necessary to store head and tail pointers indicating the present top and bottom of the FIFO queue. In the context of an Ethernet or the like computer network, it is also necessary to store head status and tail status pointers which indicate the positions it is necessary to go back to in the event a collision occurs in the transmission of the communication. Other parameters which may typically be stored include a size pointer and a shadow size pointer. The use of these pointers and the calculations which may be performed using them are well known and will not be described in detail here.

However, as can be seen, in the known arrangement illustrated in Figure 2, it is necessary to replicate the calculation device 24 with its associated registers 26 for each of the MAC devices implemented on chip 10. In the case where it is desired to implement as many as twelve MAC devices on a single chip 10 this requirement to replicate the calculating device and the associated register means there is an extremely high gate count and potentially precludes the

implementation of such device.

Figure 3 illustrates an embodiment of the present invention in which those parts which correspond to parts illustrated in Figures 1 and 2 have the same reference numerals.

5 It will be seen that, in Figure 3, there is no longer provided a separate calculation means 24 for each of the MACs provided on chip 10. Rather, there is a single calculation device 124 provided which services all of the MACs on the chip. This is of course possible because only one of the MACs has access to the SRAM 40 at any one time, and therefore only one of the MACs is requiring
10 calculations to be conducted at any time.

Associated with calculation means 124 is memory means 130 which stores the values of the parameters for each of the MACs on the chip and supplies appropriate values to calculation means 124 according to the MAC which is currently enabled by SRAM arbiter 44. This arrangement provides a considerable
15 simplification and reduction in the number of gates required in the implementation of the multiple MAC arrangement as compared to the arrangement illustrated in Figure 2.

In a particularly preferred arrangement, memory 130 is divided in to a plurality of areas 131. Each of these areas stores a particular type of
20 parameter which is used in the calculation conducted by calculating means 124 and each area 131 stores the value of that parameter for each of the MAC devices. Within each area 131 the values of the parameter are indexed according to which MAC it relates to and therefore the indication provided by SRAM arbiter 44 of which MAC currently has access to SRAM 40 can simply be used as an index
25 pointer into each area 131 to enable the provision of the relevant parameters to calculating means 124. This is a particularly simple and effective way to implement the storage of various parameters in storage means 130.

Alternatively, each of areas 131 may store all of the parameters for one of the MAC devices, with the indication provided by SRAM arbiter 44

identifying which set of parameters should be provided to calculating means 24.

It will therefore be seen that the present invention provides considerable advantages in terms of ease of implementation and cost of the final product in situations where a plurality of user devices share access on a time division basis to a shared resource, and has particular advantages in the field of the
5 implementation of multiple MAC devices on a single chip.

CLAIMS:

1. A communication device for a computer network in which a plurality of network devices are interconnected and enabled to send communications to each other, the device comprising:

a plurality of ports via which communications may be received and re-transmitted;

a communications core arranged to re-transmit received communications via one or more of said ports according to a defined functionality;

transmit means associated with each of said ports each arranged to receive from said communications core communications to be transmitted via the respective port and to transmit said communications via said respective port;

buffer memory means associated with said transmit means and arranged to have temporarily stored therein communications received by said transmit means prior to their transmission, each port having associated therewith a plurality of pointers related to the corresponding use by the transmit means of said buffer memory means;

buffer memory access means arranged to control access to said buffer memory means by said transmit means for each of said plurality of ports according to the corresponding plurality of pointers; and

pointer memory means arranged to store said pointers for each of said ports and to provide to said buffer memory access means said pointers associated with one of said ports when said buffer access means permits access to said buffer memory means by said transmit means for said one of said ports.

2. A communication device according to claim 1 in which a separate transmit means is provided for each of said plurality of ports.

3. A communication device according to claim 1 in which said buffer

memory access means comprises calculation means arranged to perform calculations based upon the values of said pointers in association with access to said buffer memory means.



5 4. A communication device according to claim 1, 2 or 3 in which said pointer memory means comprises a plurality of storage areas each storing corresponding pointers associated with each of the plurality of ports indexed according to said ports.

10 5. A method for controlling access to a shared resource by a plurality of users of said resource, each of said users having associated therewith a plurality of parameters relating to that user's use of said resource, the method comprising:
 storing all of said parameters associated with said plurality of users;
 providing access control means arranged to control access to said
15 said resource by each of said users in accordance with said parameters associated with said user;
 providing the parameters associated with each user to said access control means at times according to the times at which each user has access to said resource.

20

6. A method according to claim 5 in which said access control means comprises calculation means arranged to perform calculations on the basis of said parameters.

25 7. A method according to claim 5 or 6 in which said shared resource is a buffer memory and said parameters are memory pointers associated with each user's portion of said memory.



INTERNATIONAL SEARCH REPORT

International Application No
PCT/GB 98/03575

A. CLASSIFICATION OF SUBJECT MATTER

IPC 6 H04L12/44

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X A	US 5 633 865 A (SHORT PHILLIP) 27 May 1997 see column 2, line 52 - column 9, line 39 ----	1-3,5-7 4
X A	WO 97 31461 A (ZEITNET INC) 28 August 1997 see page 12, line 20 - page 17, line 22 see page 20, line 4 - page 23, line 8 ----	1-3,5-7 4
P,X A	EP 0 854 608 A (COMPAQ COMPUTER CORP) 22 July 1998 see page 2, line 52 - page 3, line 54 -----	1-3,5-7 4

☐ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "&" document member of the same patent family

Date of the actual completion of the international search

24 March 1999

Date of mailing of the international search report

01/04/1999

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

RAMIREZ DE AREL..., F

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/GB 98/03575

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5633865 A	27-05-1997	NONE	
WO 9731461 A	28-08-1997	US 5724358 A AU 2054397 A	03-03-1998 10-09-1997
EP 0854608 A	22-07-1998	JP 10215282 A	11-08-1998